This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (Currently amended) A method for scheduling ~~at least one resource~~ operations of a coprocessor in a computing system having a processor and a coprocessor, comprising:

receiving by a kernel mode driver a command buffer from a user mode driver~~,~~ ~~wherein the command buffer is formulated based on a request for coprocessor resources~~;

generating by ~~a~~ the kernel mode driver at least one direct memory access (DMA) buffer based on said command buffer ~~and~~;

generating by the kernel mode driver a resource list ~~corresponding to~~ comprising at least one graphics data memory resource associated with processing said at least one DMA buffer by said coprocessor, wherein ~~a set of at least one DMA buffer is maintained for each client context~~ said resource list includes an expected current context state when said DMA buffer is submitted for processing by said coprocessor;

~~analyzing said resource list, and based on said analyzing,~~ generating a paging DMA buffer including at least one instruction ~~that instructs~~ for the coprocessor ~~with respect~~ to move the at least one graphics data memory resource to an assigned location; and

submitting said at least one DMA buffer to said coprocessor for processing of DMA buffer contents.

2. (Original) A method according to claim 1, wherein the set of at least one DMA buffer is a software queue maintained for each application submitting requests for coprocessor resources.

3. (Original) A method according to claim 1, further comprising transmitting said at least one instruction to the coprocessor.

4. (Original) A method according to claim 1, wherein said submitting includes submitting said paging buffer to the coprocessor.

5. (Currently amended) A method according to claim 1, wherein said generating by the kernel mode driver at least one DMA buffer includes assigning priority to elements in the at least one DMA buffer.

6. (Original) A method according to claim 1, wherein said paging buffer includes at least one command that requests the coprocessor to stop processing.

7. (Currently amended) A method according to claim 6, wherein said at least one command is a fence that generates an interrupt and stalls the [[GPU]] coprocessor.

8. (Original) A method according to claim 1, further comprising patching at least one reference to the at least one memory resource with actual memory resource data.

9. (Original) A method according to claim 1, whereby said coprocessor processes said at least one DMA buffer in a different order than the order in which said at least one DMA buffer is generated.

10. (Currently Amended) A method according to claim 1, wherein said coprocessor is busy processing another DMA buffer ~~during at least one of said receiving, generating, analyzing and submitting~~ while the operations of claim 1 are performed.

11. (Original) A method according to claim 1, further comprising:

interrupting the processing of a DMA buffer submitted to the coprocessor according to said submitting; and

resuming processing of the interrupted DMA buffer.

12. (Original) A method according to claim 1, wherein said analyzing includes determining whether said at least one DMA buffer implicates a threshold amount of memory resources.

13. (Original) A method according to claim 12, wherein if said at least one DMA buffer implicates more than the threshold amount, splitting said at least one DMA buffer.

14. (Original) A method according to claim 1, wherein said at least one instruction includes at least one of an evict, page in and relocate instruction.

15. (Currently amended) A method according to claim 1, wherein said method is performed pursuant to operations of an [[An]] application programming interface ~~comprising~~ associated with computer executable modules including ~~having~~ computer executable instructions ~~for carrying out the method of claim 1~~ stored on a computer readable storage medium.

16-17. (Canceled)

18. (Currently amended) A computing device ~~having a scheduling component for scheduling at least one resource of a coprocessor in a computing system having a processor and a coprocessor~~, comprising:

a processor;

a coprocessor; and

a computer readable storage media coupled to the processor and the coprocessor, the computer readable storage media having stored computer-executable components comprising:

an input interface component for receiving ~~by the scheduling component~~ a command buffer from a user mode driver~~, wherein the command buffer is formulated based on a request for coprocessor resources from a particular client context~~;

a translator component that generates at least one direct memory access (DMA) buffer based on said command buffer [[and]];

a kernel mode driver component that generates a resource list ~~corresponding to~~ comprising at least one graphics data memory resource associated with processing said at least one DMA buffer by said coprocessor, wherein ~~a set of at least one DMA buffer is~~

~~maintained for each client context~~ said resource list includes an expected current context state when said DMA buffer is submitted for processing by said coprocessor;

a ~~memory~~ scheduler component for ~~analyzing said resource list, and based on said analyzing,~~ generating a paging DMA buffer including at least one instruction ~~that instructs~~ for the coprocessor ~~with respect~~ to move the at least one graphics data memory resource to an assigned location; and

an output interface component for submitting said at least one DMA buffer to said coprocessor for processing of DMA buffer contents.

19. (Currently amended) A computing device according to claim 18, wherein the set of at least one DMA buffer [[is]] includes a software queue maintained for each application submitting requests for coprocessor resources.

20. (Currently amended) A computing device according to claim 18, wherein said at least one instruction is transmitted to the coprocessor via the output interface component.

21. (Currently amended) A computing device according to claim 18, wherein said paging buffer is submitted to the coprocessor via the output interface component in conjunction with submitting said at least one DMA buffer.

22. (Original) A computing device according to claim 18, wherein priority is assigned to elements in the at least one DMA buffer.

23. (Currently amended) A computing device according to claim 18, wherein said coprocessor is configured to stop processing when requested by said paging DMA buffer which includes at least one command that requests the coprocessor to stop processing.

24. (Currently amended) A computing device according to claim 23, wherein said at least one command is a fence that generates an interrupt and stalls the [[GPU]] coprocessor.

25. (Currently amended) A computing device according to claim 18, wherein at least one reference to the at least one graphics data memory resource is patched with actual memory resource data prior to submission of the at least one DMA buffer.

26. (Original) A computing device according to claim 18, whereby said coprocessor processes said at least one DMA buffer in a different order than the order in which said at least one DMA buffer is generated.

27. (Original) A computing device according to claim 18, wherein said coprocessor is busy processing another DMA buffer while said at least one DMA buffer is being generated.

28. (Original) A computing device according to claim 18, wherein processing of a DMA buffer submitted to the coprocessor is interrupted and then processing of the interrupted DMA buffer is resumed.

29. (Currently amended) A computing device according to claim 18, wherein said ~~memory~~ scheduler component determines whether said at least one DMA buffer implicates a threshold amount of memory resources.

30. (Original) A computing device according to claim 29, wherein if said at least one DMA buffer implicates more than the threshold amount, said at least one DMA buffer is split.

31. (Original) A computing device according to claim 18, wherein said at least one instruction includes at least one of an evict, page in and relocate instruction.

32. (Currently amended) ~~A method for scheduling at least one resource~~ A computing device, including a processor and a coprocessor, configured to schedule

operations of [[a]] the coprocessor ~~in a computing system having a processor and a coprocessor~~, the computing device comprising:

means for receiving by a kernel mode driver a command buffer from a user mode driver~~, wherein the command buffer is formulated based on a request for coprocessor resources~~;

~~first~~ means for generating by the kernel mode driver at least one DMA buffer based on said command buffer ~~and~~;

means for generating by the kernel mode driver a resource list ~~corresponding to~~ comprising at least one graphics data memory resource associated with processing said at least one direct memory access (DMA) buffer by said coprocessor, wherein ~~a set of at least one DMA buffer is maintained for each client context~~ said resource list includes an expected current context state when said DMA buffer is submitted for processing by said coprocessor;

~~means for analyzing said resource list,~~

~~second~~ means for generating a paging DMA buffer~~, based on said analyzing by said means for analyzing,~~ including at least one instruction ~~that instructs~~ for the coprocessor ~~with respect~~ to move the at least one memory resource to an assigned location; and

means for submitting said at least one DMA buffer to said coprocessor for processing of DMA buffer contents.

33. (Original) A computing device according to claim 32, wherein the set of at least one DMA buffer is a software queue maintained for each application submitting requests for coprocessor resources.

34. (Original) A computing device according to claim 32, further comprising means for transmitting said at least one instruction to the coprocessor.

35. (Original) A computing device according to claim 32, wherein said means for submitting includes means for submitting said paging buffer to the coprocessor.

36. (Currently Amended) A computing device according to claim 32, wherein said ~~first~~ means for generating by the kernel mode driver at least one DMA buffer includes means for assigning priority to elements in the at least one DMA buffer.

37. (Currently amended) A computing device according to claim 32, wherein said coprocessor is configured to stop processing when requested by said paging buffer which includes at least one command that requests the coprocessor to stop processing.

38. (Currently amended) A computing device according to claim 37, wherein said at least one command is a fence that generates an interrupt and stalls the [[GPU]] coprocessor.

39. (Original) A computing device according to claim 32, further comprising means for patching at least one reference to the at least one memory resource with actual memory resource data.

40. (Original) A computing device according to claim 32, whereby said coprocessor processes said at least one DMA buffer in a different order than the order in which said at least one DMA buffer is generated by said first means for generating.

41. (Original) A computing device according to claim 32, wherein said coprocessor is busy processing another DMA buffer during operation of at least one of said means for receiving, first and second means for generating, means for analyzing and means for submitting.

42. (Original) A computing device according to claim 32, further comprising:

means for interrupting the processing of a DMA buffer submitted to the coprocessor according to said submitting; and

means for resuming processing of the interrupted DMA buffer.

43. (Original) A computing device according to claim 32, wherein said means for analyzing includes means for determining whether said at least one DMA buffer implicates a threshold amount of memory resources.

44. (Original) A computing device according to claim 43, further comprising means for splitting a DMA buffer, wherein if said at least one DMA buffer implicates more than the threshold amount, said means for splitting splits said at least one DMA buffer.

45. (Original) A computing device according to claim 32, wherein said at least one instruction includes at least one of an evict, page in and relocate instruction.